Reinforcement Learning-based Traffic Control to Optimize Energy Usage and Throughput (HPC4Mobility)

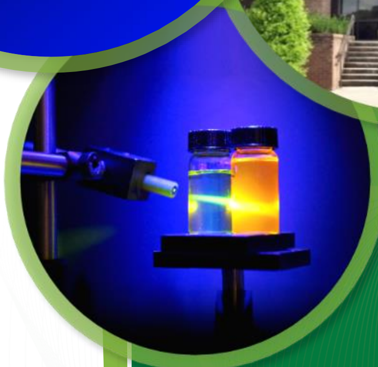

Thomas P Karnowski
Email: karnowskitp@ornl.gov
Phone: 865-574-5732

Oak Ridge National Laboratory
National Transportation Research Center

2019 U.S. DOE Vehicle Technologies Office
Annual Merit Review

June 11, 2019 **Project ID:EEMS036**

HPC4 Mobility

This presentation does not contain any proprietary, confidential, or otherwise restricted information

ORNL is managed by UT-Battelle for the US Department of Energy

Oak Ridge National Laboratory | National Transportation Research Center

Overview

Timeline

- Project start date: Feb 22, 2018
- Project end date: April 30, 2019
- Percent complete: 90%

Budget

- Total project funding:
  - DOE share: $250K
  - Industry share: $60K
- FY 2017 Funding: $NA
- FY 2018 Funding: $250K / $60K
- FY 2019 Funding: (Carryover from FY 2018)

Barriers

- Continuously adapting intelligent transportation infrastructure operations in metropolitan corridors can require observation beyond human capacity
- Optimizing traffic in urban areas may introduce hundreds of thousands of vehicles and traffic systems into a single model that requires HPC

Partners

- GRIDSMART (Knoxville, TN) CRADA collaborator

OAK RIDGE National Laboratory | NATIONAL TRANSPORTATION RESEARCH CENTER

Relevance

- **This project will create a technology solution to increase mobility energy productivity*** through energy-efficient adaptive traffic control
  - Use existing camera technology (GRIDSMART), computer vision, and machine learning to measure the energy consumption characteristics of the traffic environment
  - Use data science and High Performance Computing (HPC) to develop improved traffic light timing to allow more energy efficient control with reinforcement learning
  - **Idling costs ~6 billion gallons of fuel annually**
- Objectives this period (through March 2019)
  - Completed data collection with analysis and machine learning to detect and predict fuel consumption using GRIDSMART cameras
  - Performed simulations using reinforcement learning (RL) with visual sensing model based on GRIDSMART camera data

GRIDSMART horizon-to-horizon camera view
Copyright © GRIDSMART. Used by permission

*https://www.energy.gov/eere/vehicles/energy-efficient-mobility-systems

https://www.energy.gov/eere/vehicles/articles/new-initiatives-will-use-supercomputers-improve-transportation-energy

Approach / Strategy : Data Focus

- Use real-world imaging data and vehicle dynamic data from naturalistic driving studies to estimate fuel consumption of imaged vehicles
- Build a training set of ground-level images and corresponding GRIDSMART images to perform vehicle classification

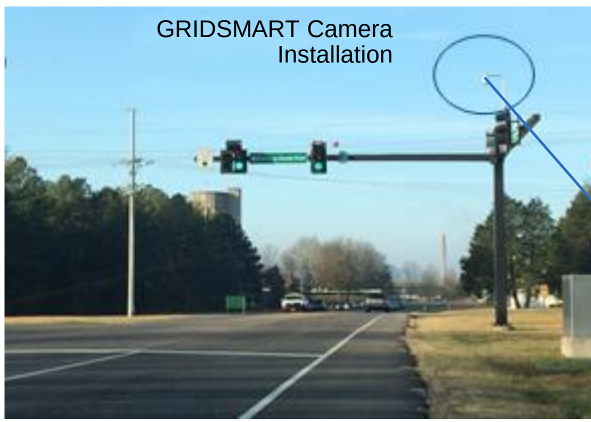

Simultaneous ground detections and GRIDSMART imagery will be correlated to develop a training set for ***vehicle classification by GRIDSMART cameras***

Ground sensor capture
(Leverages existing vehicle detection work)

Ground Truth Make Model

GRIDSMART capture tagged with make/model

OAK RIDGE National Laboratory | NATIONAL TRANSPORTATION RESEARCH CENTER

Approach / Strategy : Simulation Focus

- Simulations will be performed using historic data for traffic and GRIDSMART sensing with new computational intelligence capabilities to demonstrate adaptive signal control across large-scale urban areas
- Simulations will be used to train and test reinforcement learning algorithms that optimize energy efficiency and throughput
- Started with single intersections with a goal of larger grids

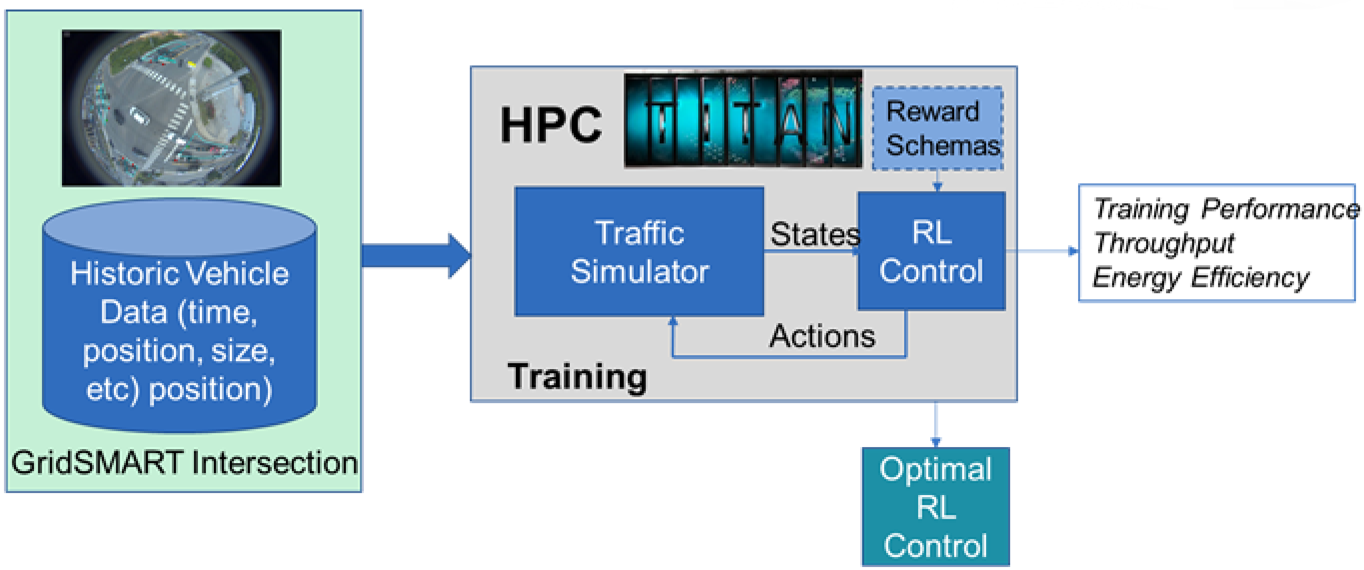

OAK RIDGE National Laboratory | NATIONAL TRANSPORTATION RESEARCH CENTER

Milestones

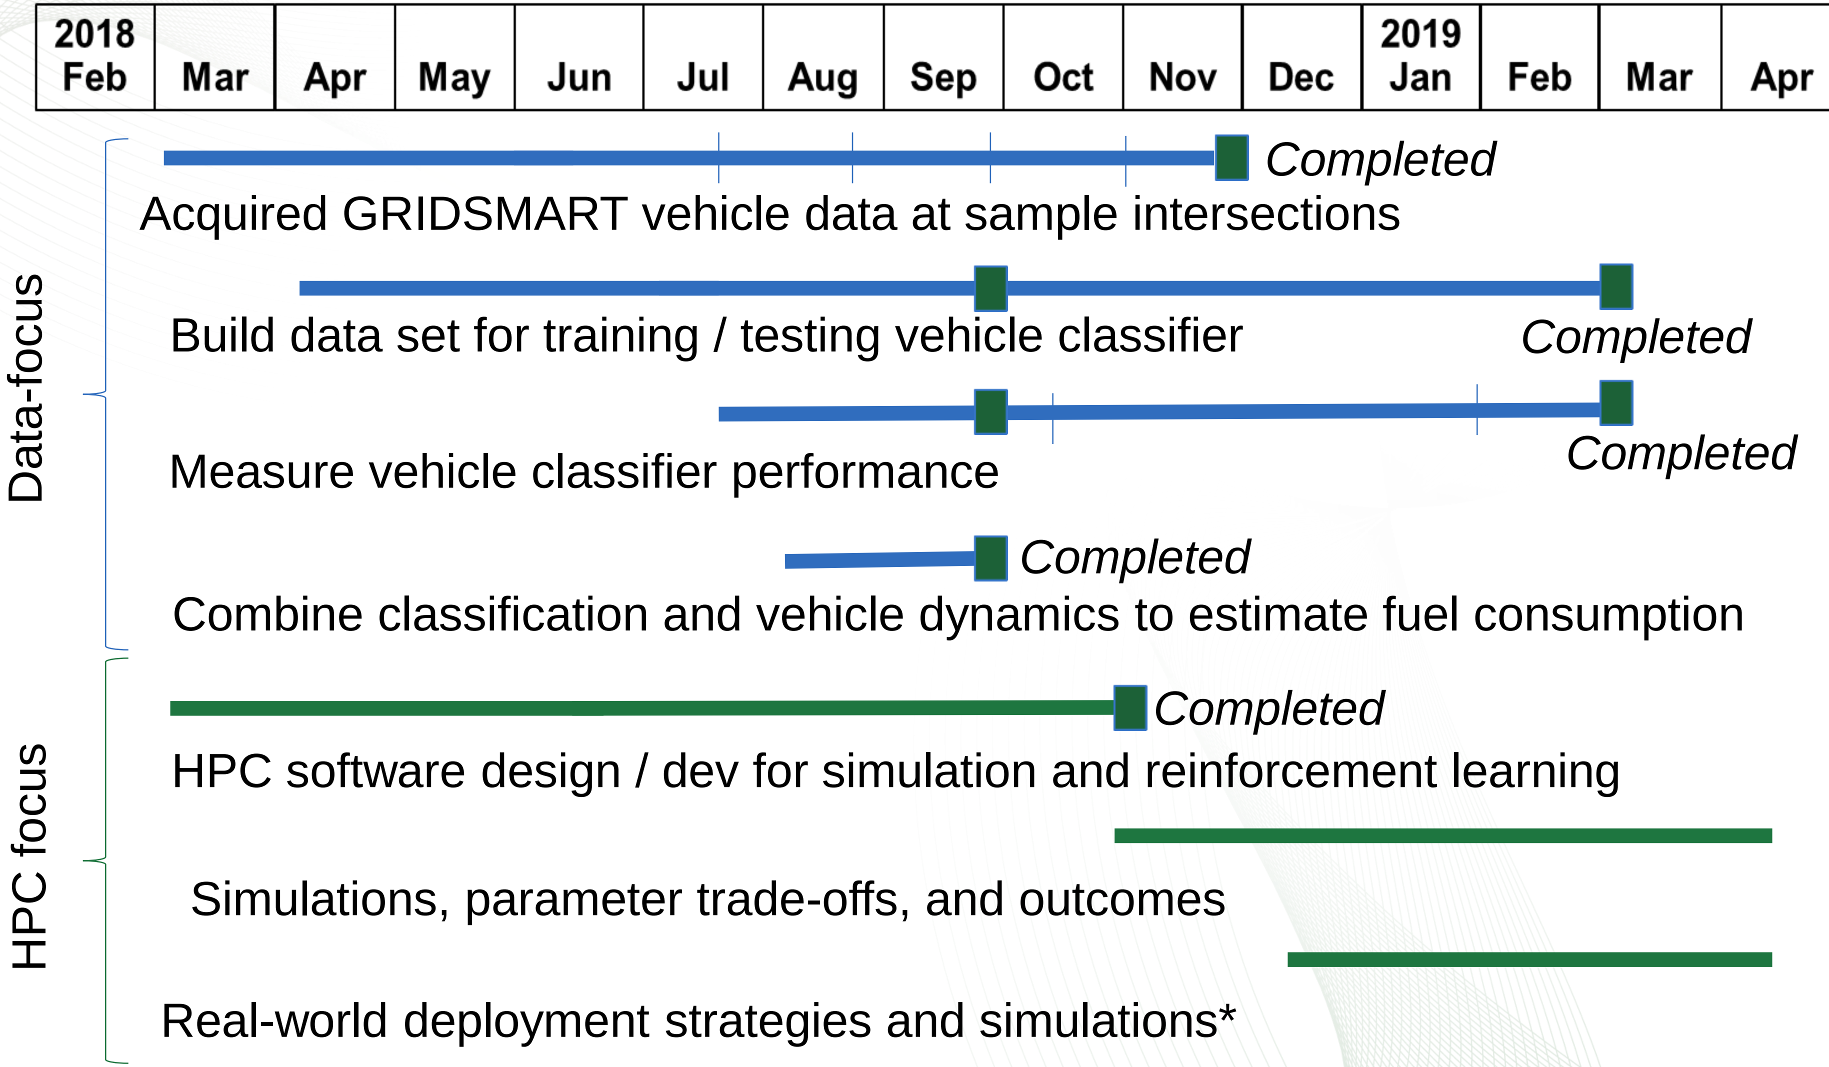
2018 Feb
Mar
Apr
May
Jun
Jul
Aug
Sep
Oct
Nov
Dec
2019 Jan
Feb
Mar
Apr
Data-focus
Completed
Acquired GRIDSMART vehicle data at sample intersections
Build data set for training / testing vehicle classifier
Completed
Completed
Measure vehicle classifier performance
Completed
Combine classification and vehicle dynamics to estimate fuel consumption
HPC focus
Completed
HPC software design / dev for simulation and reinforcement learning
Simulations, parameter trade-offs, and outcomes
Real-world deployment strategies and simulations*

Technical Accomplishments and Progress: Data Focus

- Improved real-world collection capability by time-syncing GRIDSMART cameras at ORNL with NTP time server
  - This improved our synchronization from many minutes to seconds, but we still required a considerable amount of “data hygiene”

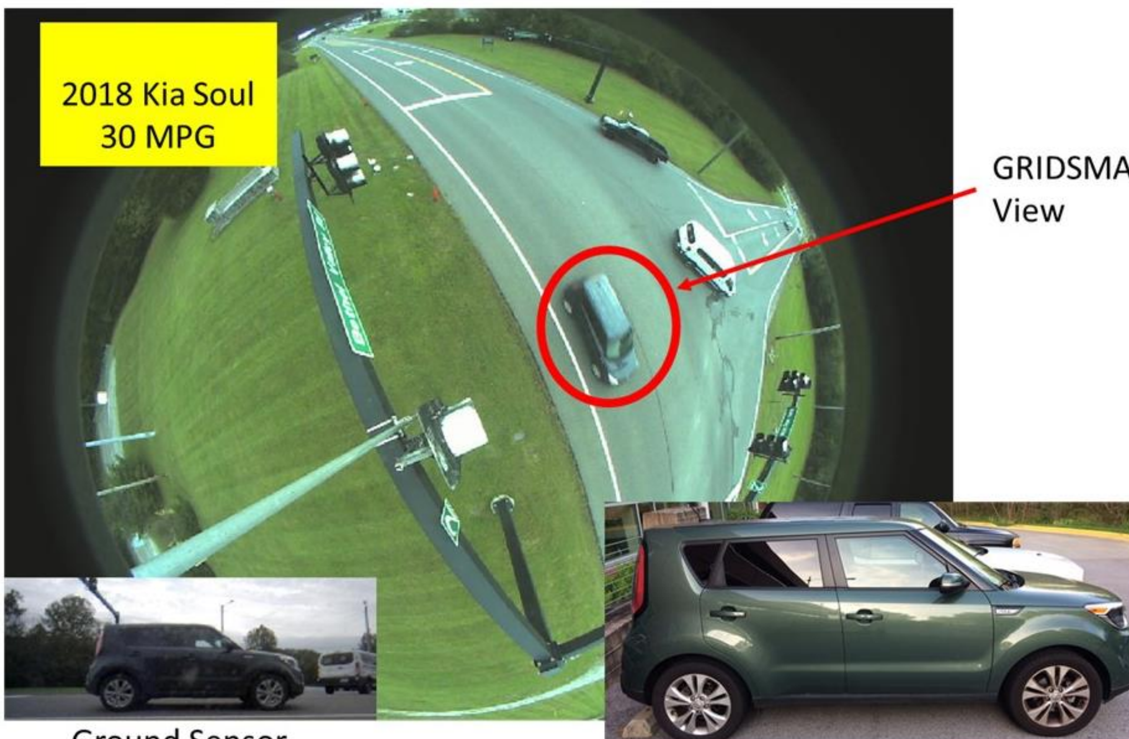

GRIDSMART View

Ground Sensor View

Photograph

Technical Accomplishments and Progress: Data Focus

- We have built a dataset of over 8000 vehicle images collected from ground image and GRIDSMART camera with machine segmentation and vehicle classes
  - We plan to make this publicly available **pending GRIDSMART and ORNL approval**

- Computer vision segmentation
- Vehicle groupings covering ~400 classifications
- Time-of-day for collection
- Traffic control experts suggested set should focus on the vehicle position with highest resolution

Technical Accomplishments and Progress: Data Focus

- Used pre-trained convolutional neural network (Krizhevsky) modified to classify vehicles by make & model and body style using data set from Gebru
- Used classifications to estimate fuel consumption “visually” and studied degradation effects of distance to vehicle
- Conclusion: Vehicle mpg classification is coarse, and only better than using the mean consumption at close range

Scaling	Est. Range	Accuracy		RMS MPG Error	
		Make Model	Body Style	Make Model	Body Style
1 x	NA	39%	74%	3.1	5.5
1/2 x	0 m	33%	69%	3.5	5.3
1/4 x	20 m	16%	60%	5.1	5.4
1/8 x	40 m	3%	41%	6.7	5.5
1/16 x	60 m	1%	14%	10.0	5.4

Krizhevsky,A. et al. "Imagenet classification with deep convolutional neural networks." *Advances in neural information processing systems*. 2012.

Gebru, T., et al. "Fine-Grained Car Detection for Visual Census Estimation." *AAAI*. Vol. 2. No. 5. 2017.

Technical Accomplishments and Progress: Data Set Utility

- The ORNL / GRIDSMART data set can be used to illustrate the utility of "bounding boxes" to detect "fuel consumption" visually
- Confusion due to bounding boxes is significantly corrected using the data set to train a vehicle classifier, with the caveat that there ARE a limited number of vehicles in the data set

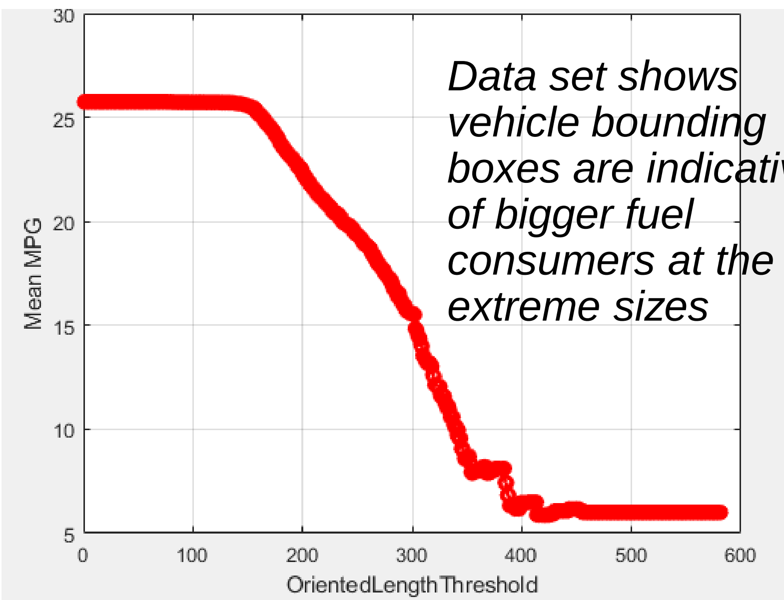

Examples of vehicles with large bounding boxes; classifier based on ResNet101 separates these classes with roughly 85% accuracy

OAK RIDGE National Laboratory | NATIONAL TRANSPORTATION RESEARCH CENTER

Technical Accomplishments and Progress: Simulation Focus

- Ported Simulation of Urban Mobility (SUMO), an open source traffic simulator, to ORNL's TITAN HPC platform and performed basic experiments
- Added simulations of GRIDSMART-like visual sensing, focusing on distinguishing large fuel consumers (i.e., trucks)
- Incorporated Deep Q-Learning reinforcement learning network (DQN) for control to compare four methods:
  - Baseline SUMO control policy
  - Heuristic vision-based policy
  - DQN reinforcement learning
  - DQNAdjacent (uses adjacent intersection visual information)

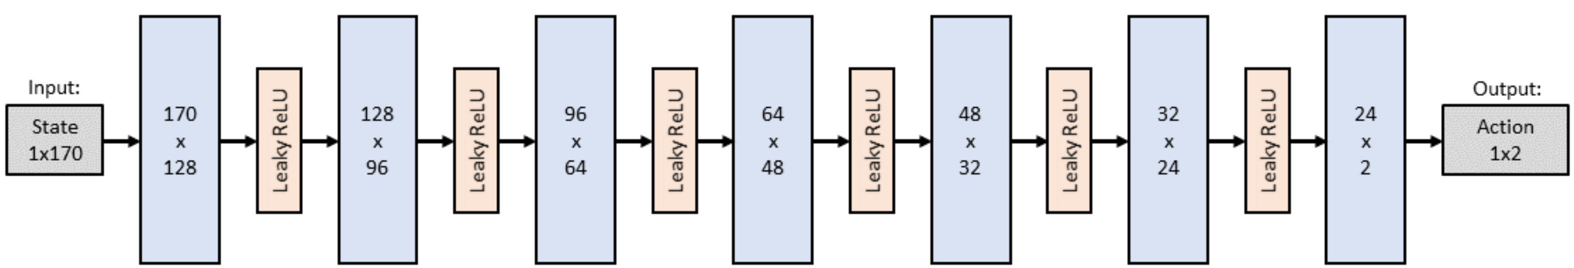

Deep network topology for control

Technical Accomplishments and Progress: Simulation Focus

- Compared standard control with three vision-based methods (heuristic policy, DQN, and DQN with input from adjacent traffic cameras)
- Preliminary results show that distinguishing large vehicles from conventional vehicles is effective for control methods

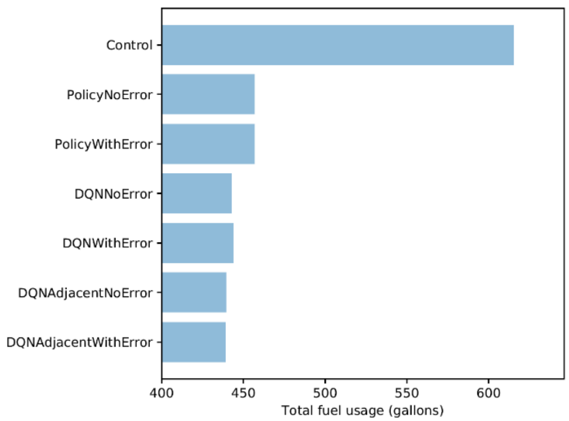

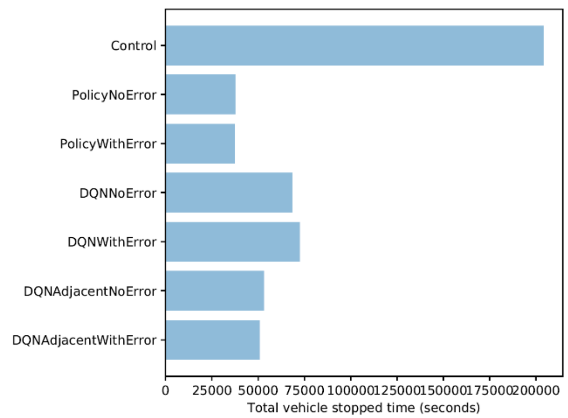

Control

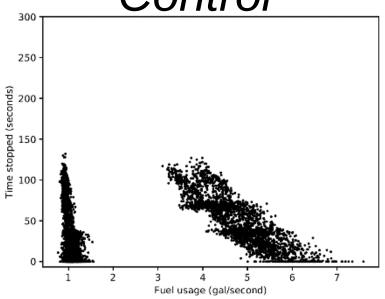

Policy

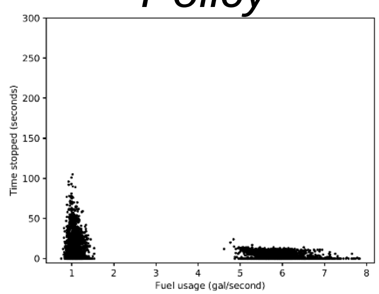

DQN

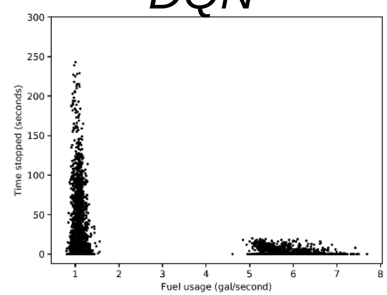

DQNAdjacent

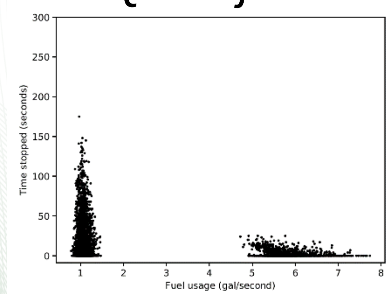

We can teach traffic light controls how to reduce fuel consumption

OAK RIDGE National Laboratory | NATIONAL TRANSPORTATION RESEARCH CENTER

Technical Accomplishments and Progress: Simulation Focus

- Sparse traffic simulations have interesting explainable result when adjacent intersection information is utilized

DQNAdjacent

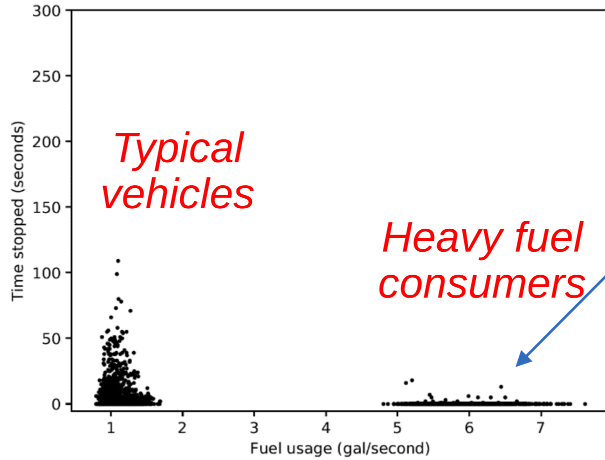

Flat line means the controller learned to ensure large trucks and buses did not stop (much)

Control

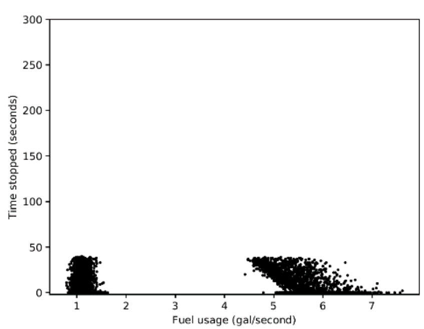

Policy

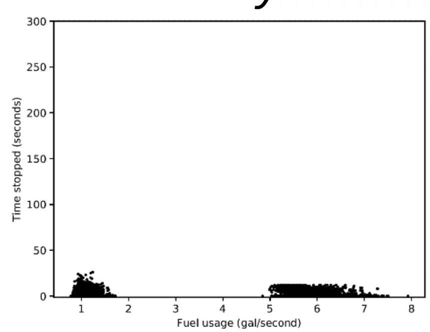

DQN

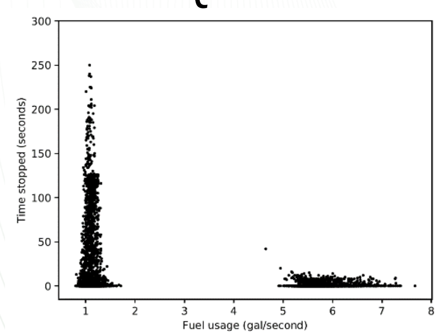

We can teach traffic light controls how to reduce fuel consumption

Technical Accomplishments and Progress: Simulation Focus

- We obtained a data set of drivers passing through intersections from the Strategic Highway Research Program 2 (SHRP2) naturalistic driving study
- Time series data includes vehicle make/model, vehicle dynamics, intersection lat/lon, engine RPM, and radar data
- Data set is publicly available provided data usage license is signed with Virginia Tech Transportation Institute (VTTI)
  - https://doi.org/10.15787/VTT1/4JOEAP
- However, since we used SUMO for simulations, we did not incorporate the data into our modeling effort

Response to Previous Year's Comments

- Comment: Would like more information on real-world deployment strategies
  - Response: Real-world deployment would require expanding beyond GRIDSMART to include traffic controller vendors/collaborators as well as a higher level of integration. We believe our work advances the feasibility of the concept but recognize that more R&D is required at this stage (a concurring statement was made by Reviewer 6)
- Comment: Go/No-Go metrics were not shown nor were there alternate approaches
  - Response: We did have a contingency plan with respect to the vehicle classification data. Reinforcement Learning has been applied to traffic problems before and therefore we believed (and still believe) it is a viable approach*, although for our experiments we added a baseline "heuristic" policy that functions as a contingency / alternate approach. Finally we believe the scope of the project is such that alternate approaches would be best reserved for future research.

**"Minimizing energy consumption from connected signalized intersections by reinforcement learning", Bin Al Islam et al, 2018, part of Urban Science Task 4.2 of DOE SMART initiative*

Partners / Collaborators

- GRIDSMART: CRADA partner
  - Supplies intersection data at image level and statistics level (vehicle size and maneuver counts)
  - Provided technical expertise on camera / network integration
  - Provided understanding of processing methods currently employed by GRIDSMART sensors including vehicle segmentation, tracking, and characteristics
  - Provided subject matter technical guidance
- Other ***future*** partners could include municipalities that utilize GRIDSMART technology:
  - Allentown PA
  - Sevierville TN
  - Chattanooga TN

GRIDSMART®
TRAFFIC CONTROL FREAKS

Remaining Challenges and Barriers

- Data focus
  - Data set required considerable data hygiene which could be improved through better timing synchronization, better ground truth (such as license plate readers) and improvements to ground camera-based acquisition
  - Plan to complete release of dataset for R&D purposes (pending ORNL and GRISMART approval)
- Simulation focus
  - SUMO was ported to ORNL TITAN HPC platform
  - Completing simulations with varied parameters on single intersection and small grid with visual detection / classification models
  - We were *not* able to take full advantage of the HPC paradigm within the project
- CRADA completion at end of April 2019 (one year project, with two-month no-cost extension) with final report

Proposed Future Research

- Data focus
  - Improved data sets with recent improvements to ground-based camera system, including better synchronization with GRIDSMART cameras, license plate readers, and better ground-truth vehicle classifications
  - Better groupings of vehicles for fuel consumption estimates
  - Larger data sets from multiple intersections, including non-ORNL campus
- Simulation focus
  - Simulations beyond this project include using MENNDL to evolve reward structure and deep-Q topology, especially for larger grids
  - Simulations could expand to real-world data and digital twins, particularly in partnership with other projects that leverage GRIDSMART technology
  - MENNDL could also be employed to create classification strategies that improve fuel consumption estimates
  - MENNDL could evolve convolutional neural networks with hardware constraints suitable for GRIDSMART deployment with appropriate hardware

Summary

- **Relevance:** Create a technology solution to increase mobility energy productivity using GRIDSMART cameras, computer vision, and reinforcement learning for adaptive signal control
- **Approach:**
  - Use historic data and imagery to create methods to estimate fuel consumption of vehicles at intersection
  - Use these estimates of camera performance in simulations with reinforcement learning to develop energy-efficient traffic control
- **Collaborations:** GRIDSMART of Knoxville, TN
- **Technical Accomplishments:**
  - Obtained ground truth data at ORNL intersection and built public* dataset
  - Specified naturalistic driving study dataset of intersection traversals
  - Adapted baseline methods for visual fuel consumption estimation
  - Ported SUMO to HPC platform
  - Tested Deep-Q reinforcement learning algorithms for visual-sensed traffic control with SUMO traffic simulations and Keras/TensorFlow machine learning toolkits
- **Future Work**
  - Improved data sets
  - Larger simulations
  - Integration into other projects leveraging GRIDSMART technology

HPC4 Mobility

Any proposed future work is subject to change based on funding levels

TECHNICAL BACKUP SLIDES

Potential Follow-on Projects

- **Improved Data Sets:** Leverage lessons learned to create better data sets at multiple intersections
- **Larger Simulations:** Advances in simulation capability are significant but more can be done with more extensive runs with larger grids, including MENNDL technology to evolve better networks and control policies
- **Integration into projects leveraging GRIDSMART technology**: The “situational awareness” offered by the visual data sets and classification can augment and enhance projects developing digital twins and data analytics
- **Testbed evaluation: ORNL could serve as a testbed for actual control work (pending appropriate collaboration partners)**

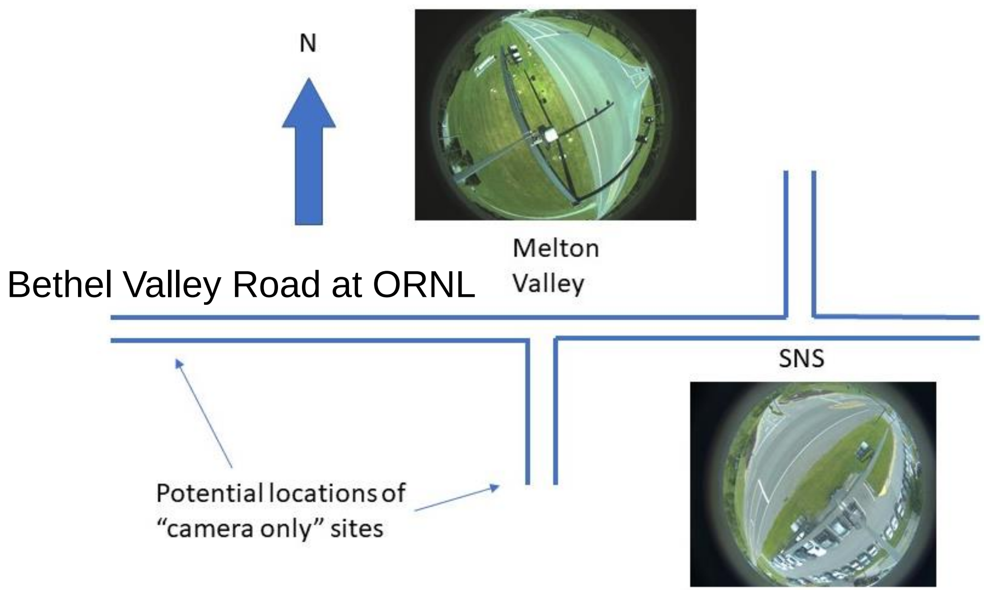

Any proposed future work is subject to change based on funding levels

Data set vehicle classification examples

- Vehicle bounding box from segmentation can be an effective discriminator between very large vehicles with high fuel consumption
- A visual classifier, trained on the dataset using features extracted from a pre-trained RESNET101 convolutional neural network (He) and a support vector machine (SVM) classifier, can separate large bounding-box vehicle classes with roughly 85% accuracy on initial testing conducted on our data set

Estimated vehicle class

Actual vehicle class

	Bus & 18 Wheeler	Express Bus	Delivery Van	Multi Axel	Passenger Vehicles*
Bus & 18Wheeler	0.78	0.0	0.17	0.06	0.0
Express Bus	0.0	0.94	0.06	0.0	0.0
Delivery Van	0.06	0.0	0.89	0.06	0.0
Multi Axel	0.06	0.0	0.22	0.72	0.0
Passenger Vehicles*	0.0	0.0	0.0	0.0	1.0

He, et al. "Deep residual learning for image recognition." *Proc. IEEE CVPR*, 2016.

**Features passenger vehicles with largest bounding boxes as potential confusers*

Data Focus: Improving Visual MPG Estimate

- ***Use High-Performance Computing to evolve a better MPG estimate!***

Multinode Evolutionary Neural Networks for Deep Learning

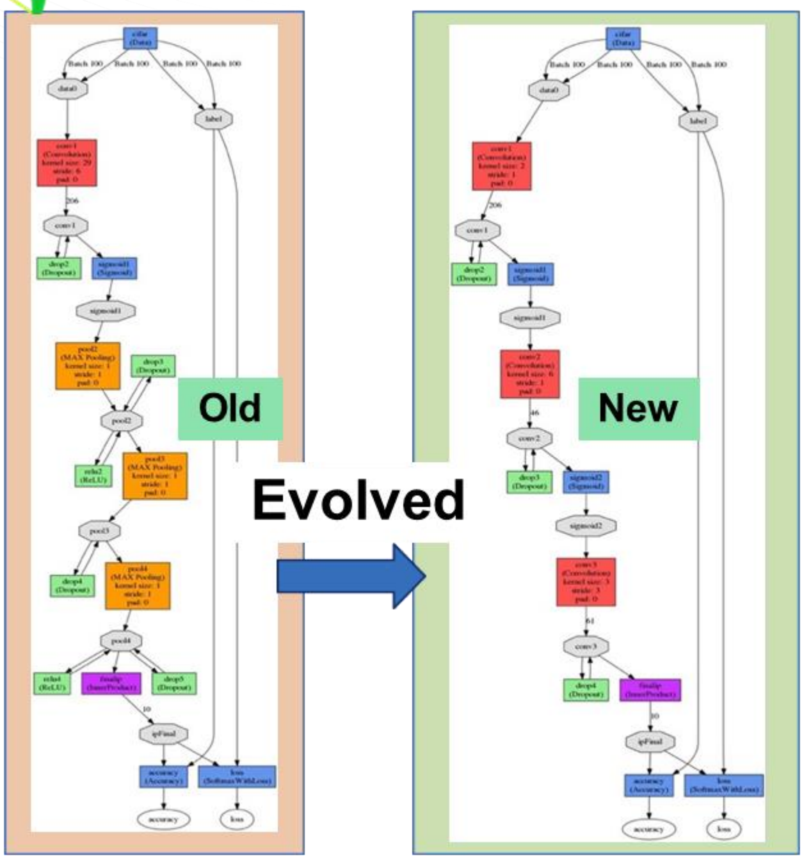

- Deep Learning must be tailored to an application domain, which can require months of manual effort with a significant amount of expertise
- Mutitnode Evolutionary Neural Networks for Deep Learning (MENNDL) automatically creates and evolves DL neural networks that best suit an application and runs on 100% of Titan and Summit
- *MENNDL creates networks that typically go beyond what a human expert would have considered*